

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants:	Bahman Zargham, et al.	Examiner:	Merilyn P. Nguyen
Serial No.:	10/013,091	Group Art Unit:	2163
Filed:	December 7, 2001	Docket No.:	200302218-1
Title:	Zero Latency Enterprise Enriched Published/Subscribed		

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is filed in response to the Final Office Action mailed March 21, 2007 and Notice of Appeal filed on July 5, 2007.

AUTHORIZATION TO DEBIT ACCOUNT

It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's deposit account no. 08-2025.

I. REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. RELATED APPEALS AND INTERFERENCES

There are no known related appeals or interferences known to appellant, the appellant's legal representative, or assignee that will directly affect or be directly affected by or have a bearing on the Appeal Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1, 2, 4-20, and 22-27 stand finally rejected. Claims 3 and 21 are objected to as being dependent upon a rejected base claim but would be allowed if rewritten in independent form including all the limitations of the base claims and any intervening claims. The rejection of claims 1, 2, 4-20, and 22-27 is appealed.

IV. STATUS OF AMENDMENTS

No amendments were made after receipt of the Final Office Action. All amendments have been entered.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R. § 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations

from the specification and drawings should be read into the corresponding claim element or that these are the sole sources in the specification supporting the claim features.

Claim 1 is directed to a method for enriched publish and subscribe in an enterprise running as a zero latency enterprise (ZLE), the enterprise experiencing a plurality of events occurring in association with business transactions conducted at a plurality of sites across the enterprise (see p. 3, line 20 – p. 4, line 20: describing publish and subscribe operations in an IT infrastructure to allow an enterprise to runs a zero latency enterprise, ZLE). The method comprises initiating, in real time, a process responsive to an event, the process including (see p. 8, line 17 – p. 9, line 2 for definitions of real time and zero latency), publishing to a central repository (Figures 1-6, #106) one or more messages prompted by that event containing information from that event, respective information from the plurality of events being aggregated in the central repository where the aggregated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise (see p. 7, lines 15-19 for definition of publish; p. 7, lines 20-23 for definition of enterprise; p. 10, lines 7-27 for definition of event; p. 11, lines 21-25 for definition of central repository; p. 7, lines 24-27 for definition of across the enterprise; message router 124 routes messages to central repository, see p. 16, line 5 – p. 17, line 6; see p. 22 for publish and subscribe functionality), updating the aggregated information with information from the published messages (see see p. 22 for publish and subscribe functionality; example of updating information in Figure 6 at p. 26, lines 22-30), creating enriched messages by enriching the messages with information from that event and/or corresponding information extracted from the central repository (see section on enriched publish and subscribe at p. 23, example, the ability of the ODS to cache data to enrich messages), and subscribing the enriched messages (see section on enriched publish and subscribe at p. 23; Figure 5 shows enterprise example for product ordering, manufacturing, and shipping).

Claim 11 is directed to a system for providing enriched publish and subscribe in an enterprise running as a zero latency enterprise (ZLE), the enterprise experiencing a plurality of events occurring in association with business transactions conducted at a plurality of sites across the enterprise (see p. 3, line 20 – p. 4, line 20: describing publish and subscribe operations in an IT infrastructure to allow an enterprise to runs a zero

latency enterprise, ZLE). The system comprises means for initiating, in real time, a process responsive to an event, the means for initiating the process including (see p. 8, line 17 – p. 9, line 2 for definitions of real time and zero latency), means for publishing to a central repository (Figures 1-6, #106) one or more messages prompted by that event containing information from that event, respective information from the plurality of events being aggregated in the central repository where the aggregated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise (see p. 7, lines 15-19 for definition of publish; p. 7, lines 20-23 for definition of enterprise; p. 10, lines 7-27 for definition of event; p. 11, lines 21-25 for definition of central repository; p. 7, lines 24-27 for definition of across the enterprise; message router 124 routes messages to central repository, see p. 16, line 5 – p. 17, line 6; see p. 22 for publish and subscribe functionality), means for updating the aggregated information with information from the published messages (see see p. 22 for publish and subscribe functionality; example of updating information in Figure 6 at p. 26, lines 22-30), means for creating enriched messages by enriching the messages with information from that event and/or corresponding information extracted from the central repository (see section on enriched publish and subscribe at p. 23, example, the ability of the ODS to cache data to enrich messages), and means for subscribing the enriched messages (see section on enriched publish and subscribe at p. 23; Figure 5 shows enterprise example for product ordering, manufacturing, and shipping).

Claim 15 is directed to a computer readable medium embodying computer program code configured to cause a computer to perform steps for providing enriched publish and subscribe in an enterprise running as a zero latency enterprise (ZLE), the enterprise experiencing a plurality of events occurring in association with business transactions conducted at a plurality of sites across the enterprise (see p. 3, line 20 – p. 4, line 20: describing publish and subscribe operations in an IT infrastructure to allow an enterprise to runs a zero latency enterprise, ZLE). The steps comprise initiating, in real time, a process responsive to an event, the process including the steps of (see p. 8, line 17 – p. 9, line 2 for definitions of real time and zero latency), publishing to a central repository (Figures 1-6, #106) one or more messages prompted by that event containing information from that event, respective information from the plurality of events being

aggregated in the central repository where the aggregated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise (see p. 7, lines 15-19 for definition of publish; p. 7, lines 20-23 for definition of enterprise; p. 10, lines 7-27 for definition of event; p. 11, lines 21-25 for definition of central repository; p. 7, lines 24-27 for definition of across the enterprise; message router 124 routes messages to central repository, see p. 16, line 5 – p. 17, line 6; see p. 22 for publish and subscribe functionality), updating the aggregated information with information from the published messages (see p. 22 for publish and subscribe functionality; example of updating information in Figure 6 at p. 26, lines 22-30), creating an enriched message by enriching the messages with information from that event and/or corresponding information extracted from the central repository (see section on enriched publish and subscribe at p. 23, example, the ability of the ODS to cache data to enrich messages), and subscribing the enriched messages (see section on enriched publish and subscribe at p. 23; Figure 5 shows enterprise example for product ordering, manufacturing, and shipping).

Claim 16 is directed to a ZLE virtual hub (#102) for enriched publish and subscribe operations associated with business transactions conducted by an enterprise running as a zero latency enterprise (ZLE) (the ZLE concept, framework, and architecture are provided in a discussion beginning on p. 12). The system being implemented in a ZLE framework and comprises one or more applications (118 or 110 in Figures 1 and 2) via which the business transactions are conducted (see p. 14, line 3-12); and an operational data store (ODS) (#102), the ODS being operatively communicative with the one or more applications such that the applications are capable to publish messages to and subscribe to messages from the ODS (see p. 13, line 7 – p. 15, line for discussion of operation data store and ZLE framework), the ODS being configured to operate as a dynamic central repository that consolidates information from across the enterprise and supports business transactional access to real time information from any of the one or more applications (see section on ODS with Cluster-aware RDBMS beginning p. 19, discussing ODS consolidates data across the enterprise), to know what particular information any one of the applications needs in order to accomplish its task, the particular information enriching messages to which the applications subscribe (see section on ODS with Cluster-aware RDBMS beginning p. 19.), and to update the

consolidated information with information from messages published by the applications (see section on ODS with Cluster-aware RDBMS beginning p. 19 discussing ODS consolidates data from across the enterprise in real time and supports transactional access to up-to-the-second data from multiple systems and applications).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-2, 5-8, 10-13, 15-16, and 19-27 are rejected under 35 USC § 102(e) as being anticipated by US publication number 2002/0013759 (Stewart).

Claims 4 and 18 are rejected under 35 USC § 103(a) as being unpatentable over US publication number 2002/0013759 (Stewart) in view of US publication number 2002/0026630 (Schmidt).

Claims 9, 14, and 17 are rejected under 35 USC § 103(a) as being unpatentable over US publication number 2002/0013759 (Stewart) in view of USPN 6,058,389 (Chandra).

VII. ARGUMENT

The rejection of claims 1, 2, 4-20, and 22-27 is improper, and Applicants respectfully request reversal of these rejections.

The claims do not stand or fall together. Instead, Applicants present separate arguments for various claims. Each of these arguments is separately argued below and presented with separate headings and sub-heading as required by 37 C.F.R.

§ 41.37(c)(1)(vii).

Claims Rejection: 35 USC § 102(e)

Claims 1-2, 5-8, 10-13, 15-16, and 19-27 are rejected under 35 USC § 102(e) as being anticipated by US publication number 2002/0013759 (Stewart). These rejections are traversed.

Each of the independent claims recites one or more elements that are not taught in Stewart. Some examples are provided below.

Claims 1, 11, and 15

As a first example, independent claims 1, 11, and 15 recite publishing messages to a central repository. Stewart does not teach this element. The Examiner argues Stewart at paragraph [0124] at lines 6-7 teaches this element. Applicants respectfully disagree.

Paragraph [0124] in Stewart teaches that an initiator initiates a flow of messages among a c-space and such messages are published to other participants. Nowhere does Stewart teach that these messages are published to a “central repository” as recited in claims 1, 11, and 15. By contrast, Stewart expressly teaches that these messages are published to other participants so the participants can join in the conversation.

Anticipation under section 102 can be found only if a single reference shows exactly what is claimed (see *Titanium Metals Corp. v. Banner*, 778 F.2d 775, 227 U.S.P.Q. 773 (Fed. Cir. 1985)). For at least these reasons, independent claims 1, 11, and 15 and their respective dependent claims are allowable over Stewart.

As another example, independent claims 1, 11, and 15 recite that the messages are prompted by an event and contain information from the event. Stewart does not teach this

element. The Examiner argues Stewart at paragraph [0124] at lines 6-7 teaches this element. Applicants respectfully disagree.

Paragraph [0124] in Stewart teaches that an initiator initiates a flow of messages among a c-space and such messages are published to other participants. Nowhere does Stewart teach that these messages are both prompted by an event and contain information from the event. In other words, Stewart never states that the messages to other participants are prompted by a certain event and that the messages contain information from that event.

For a prior art reference to anticipate under section 102, every element of the claimed invention must be identically shown in a single reference (see *In re Bond*, 910 F.2d 831, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990)). For at least these reasons, independent claims 1, 11, and 15 and their respective dependent claims are allowable over Stewart.

As yet another example, independent claims 1, 11, and 15 recite that information from the plurality of events is aggregated at a central repository. Stewart does not teach this element. The Examiner argues Stewart at paragraphs [0201], [0139], and [0150] teach this element. Applicants respectfully disagree.

Paragraph [0201] in Stewart discusses trading partners and gaining access to c-space. This section of Stewart has nothing to do with aggregating information from plural events at a central repository.

Paragraph [0139] in Stewart discusses a workflow server process engine communicating with a client and enterprise components. This section of Stewart has nothing to do with aggregating information from plural events at a central repository.

Paragraph [0150] in Stewart discusses that a c-hub stores information. This section of Stewart has nothing to do with aggregating information from plural events at a central repository when these terms are taken in context of the words of the claims.

Anticipation is established only when a single prior art reference discloses each and every element of a claimed invention united in the same way (see *RCA Corp. v. Applied Digital Data Systems, Inc.*, 730 F.2d 1440, 1444 (Fed. Cir. 1984)). For at least these reasons, independent claims 1, 11, and 15 and their respective dependent claims are allowable over Stewart.

As yet another example, independent claims 1, 11, and 15 recite that information from the plurality of events is aggregated at a central repository. The claims then recites that this information is “accessible and available for extraction across the enterprise.” Stewart does not teach this element. The Examiner argues Stewart at paragraphs [0227] and [0279] teach this element. Applicants respectfully disagree.

Paragraph [0227] in Stewart discusses evaluating an XPATH expression against an XML document. Nowhere does this paragraph states that aggregated information in a central repository is accessible and available for extraction across the enterprise.

Paragraph [0279] in Stewart discusses constructing a c-space using unified modeling language. Nowhere does this paragraph states that aggregated information in a central repository is accessible and available for extraction across the enterprise.

In order for a prior art reference to be anticipatory under 35 U.S.C. § 102 with respect to a claim, “[t]he elements must be arranged as required by the claim,” see M.P.E.P. § 2131, citing *In re Bond*, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990). For at least these reasons, independent claims 1, 11, and 15 and their respective dependent claims are allowable over Stewart.

As yet another example, independent claims 1, 11, and 15 recite updating aggregated information with information from published messages. The Examiner argues Stewart at paragraph [0280] teaches this element. Applicants respectfully disagree.

Paragraph [0280] in Stewart discusses messages as being an abstraction of information that share common behaviors. Further, when a transactor wants to communicate with someone over the network, the transactor instantiates an adapter. Nowhere does this paragraph in Stewart state updating aggregated information with information from published messages.

There can be no difference between the claimed invention and the cited reference, as viewed by a person of ordinary skill in the art (see *Scripps Clinic & Research Foundation v. Genentech Inc.*, 927 F.2d 1565, 1576 (Fed. Cir. 1991)). For at least these reasons, independent claims 1, 11, and 15 and their respective dependent claims are allowable over Stewart.

As yet another example, independent claims 1, 11, and 15 recite enriching messages with information from the event and/or corresponding information extracted

from the central repository. The Examiner argues Stewart at paragraphs [0227-0228] teach this element. Applicants respectfully disagree.

Paragraphs [0227-0228] in Stewart discuss XPATH filters, evaluating XPATH expressions against XML documents, etc. Nowhere does this section of Stewart teach enriching messages with information from the event and/or corresponding information extracted from the central repository.

It is well settled that to anticipate a claim, the reference must teach every element of the claim, see M.P.E.P. § 2131. For at least these reasons, independent claims 1, 11, and 15 and their respective dependent claims are allowable over Stewart.

Claim 16

As one example, independent claim 16 recites an operational data store (ODS) operatively communicative with applications. The claim then recites that “the applications are capable to publish messages to and subscribe messages from the ODS.” The Examiner argues Stewart at paragraphs [0089] to [0113] teach this element. Applicants respectfully disagree.

Paragraphs [0089] to [0113] in Stewart discuss operations of the collaboration hub. The collaboration hub is “responsible for routing messages between various c-enabled components” (see Stewart at paragraph [0089]). Notice that these paragraphs of Stewart discuss the functions of the **collaboration hub, not the applications**. In other words, Stewart discusses that the collaboration hub forwards messages. By contrast, claim 16 recites that the applications (not a central collaboration hub) publish messages to and subscribe messages from the hub. Stewart does not teach that applications themselves publish messages to and subscribe messages from the collaboration hub.

Anticipation under section 102 can be found only if a single reference shows exactly what is claimed (see *Titanium Metals Corp. v. Banner*, 778 F.2d 775, 227 U.S.P.Q. 773 (Fed. Cir. 1985)). For at least these reasons, independent claim 16 and its dependent claims are allowable over Stewart.

As yet another example, claim 16 recites that the ODS knows what particular information the applications need to accomplish their tasks. The Examiner argues Stewart at paragraph [0198] teaches this element. Applicants respectfully disagree.

Paragraph [0198] in Stewart discusses that the c-space owner defines the collaboration process or means in which transactions are conducted. This section has nothing whatsoever to do with the collaboration hub knowing what particular information the applications need to accomplish their tasks.

For a prior art reference to anticipate under section 102, every element of the claimed invention must be identically shown in a single reference (see *In re Bond*, 910 F.2d 831, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990)). For at least these reasons, independent claim 16 and its dependent claims are allowable over Stewart.

As yet another example, claim 16 recites that the ODS updates consolidated information with information from messages published by applications. The Examiner argues Stewart at paragraphs [0202], [0279], and [0280] teach this element. Applicants respectfully disagree.

Paragraph [0202] in Stewart teaches c-space owners can add and remove trading partners. This section of Stewart has nothing whatsoever to do with the collaborative hub updating consolidated information with information from messages published by applications.

Paragraph [0279] in Stewart teaches constructing a c-space with unified modeling language. This section of Stewart has nothing whatsoever to do with the collaborative hub updating consolidated information with information from messages published by applications.

Paragraph [0280] in Stewart teaches message types and their behaviors. This section of Stewart has nothing whatsoever to do with the collaborative hub updating consolidated information with information from messages published by applications.

In order for a prior art reference to be anticipatory under 35 U.S.C. § 102 with respect to a claim, “[t]he elements must be arranged as required by the claim,” see M.P.E.P. § 2131, citing *In re Bond*, 15 U.S.P.Q.2d 1566 (Fed. Cir. 1990). For at least these reasons, independent claim 16 and its dependent claims are allowable over Stewart.

Claims Rejection: 35 USC § 103(a)

Claims 4 and 18 are rejected under 35 USC § 103(a) as being unpatentable over US publication number 2002/0013759 (Stewart) in view of US publication number 2002/0026630 (Schmidt). These rejections are traversed.

As noted above, Stewart fails to teach or suggest all the elements of independent claims 1 and 16. Schmidt fails to cure these deficiencies. Thus, for at least the reasons given with respect to independent claims 1 and 16, respective dependent claims 4 and 18 are allowable over Stewart in view of Schmidt.

Claims Rejection: 35 USC § 103(a)

Claims 9, 14, and 17 are rejected under 35 USC § 103(a) as being unpatentable over US publication number 2002/0013759 (Stewart) in view of USPN 6,058,389 (Chandra). These rejections are traversed.

As noted above, Stewart fails to teach or suggest all the elements of independent claims 1, 11 and 16. Chandra fails to cure these deficiencies. Thus, for at least the reasons given with respect to independent claims 1, 11, and 16, respective dependent claims 9, 14, and 17 are allowable over Stewart in view of Chandra.

CONCLUSION

In view of the above, Applicants respectfully request the Board of Appeals to reverse the Examiner's rejection of all pending claims.

Any inquiry regarding this Amendment and Response should be directed to Philip S. Lyren at Telephone No. 832-236-5529. In addition, all correspondence should continue to be directed to the following address:

Hewlett-Packard Company
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Respectfully submitted,

/Philip S. Lyren #40,709/

Philip S. Lyren
Reg. No. 40,709
Ph: 832-236-5529

VIII. Claims Appendix

1. (Previously Presented) A method for enriched publish and subscribe in an enterprise running as a zero latency enterprise (ZLE), the enterprise experiencing a plurality of events occurring in association with business transactions conducted at a plurality of sites across the enterprise, the method comprising:
 - initiating, in real time, a process responsive to an event, the process including
 - publishing to a central repository one or more messages prompted by that event containing information from that event, respective information from the plurality of events being aggregated in the central repository where the aggregated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise,
 - updating the aggregated information with information from the published messages,
 - creating enriched messages by enriching the messages with information from that event and/or corresponding information extracted from the central repository, and
 - subscribing the enriched messages.
2. (Original) A method as in claim 1, wherein the central repository operates as an information broker between applications such that applications publish messages to the central repository and subscribe to messages from the central, rather than exchange request-response messages directly with each other.
3. (Original) A method as in claim 2, wherein for a particular number (N) of applications, a combined number of the published and subscribed messages can be reduced from twice that particular number (2N) to a total number of 4 or 5 messages.
4. (Original) A method as in claim 1, wherein the central repository provides a coherent view, in real time, of the aggregated information from across the enterprise, the process being founded on the coherent view of the aggregated information.

5. (Original) A method as in claim 2, further comprising:
predefining a schema for each of the applications, the schema identifying which of the plurality of events and types of data changes its respective application is interested in, the schema further identifying any information its respective application needs for performing tasks related to such events; and
storing each schema in the central repository for later use in enriching the new messages.
6. (Original) A method as in claim 2, wherein the applications cause the updating of aggregated information at the central repository upon a change of information in their environment.
7. (Original) A method as in claim 1, wherein the published messages and subscribed messages are formatted in XML.
8. (Original) A method as in claim 1, wherein the central repository is based on a database which can be updated with information from new events while being queried and which can send the enriched messages to multiple subscribers, thereby leveraging an innate parallelism, scalability and reliability of the database.
9. (Original) A method as in claim 1, wherein the central repository includes relational database management functionality that caches and queues the published and subscribed messages.
10. (Original) A method as in claim 1, wherein the enriched new messages to which an application subscribes can include extracted information that was previously published to the central repository by other one or more applications.

11. (Previously Presented) A system for providing enriched publish and subscribe in an enterprise running as a zero latency enterprise (ZLE), the enterprise experiencing a plurality of events occurring in association with business transactions conducted at a plurality of sites across the enterprise, the system comprising:

means for initiating, in real time, a process responsive to an event, the means for initiating the process including

means for publishing to a central repository one or more messages prompted by that event containing information from that event, respective information from the plurality of events being aggregated in the central repository where the aggregated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise,

means for updating the aggregated information with information from the published messages,

means for creating enriched messages by enriching the messages with information from that event and/or corresponding information extracted from the central repository, and

means for subscribing the enriched messages.

12. (Original) A system as in claim 11, wherein the central repository operates as an information broker between applications such that applications publish messages to the central repository and subscribe to messages from the central, rather than exchange request-response messages directly with each other.

13. (Original) A system as in claim 11, wherein the central repository is based on a database which can be updated with information from new events while being queried and which can send the enriched messages to multiple subscribers, thereby leveraging an innate parallelism, scalability and reliability of the database.

14. (Original) A method as in claim 11, wherein the central repository includes relational database management functionality that caches and queues the published and subscribed messages.

15. (Previously Presented) A computer readable medium embodying computer program code configured to cause a computer to perform steps for providing enriched publish and subscribe in an enterprise running as a zero latency enterprise (ZLE), the enterprise experiencing a plurality of events occurring in association with business transactions conducted at a plurality of sites across the enterprise, the steps comprising:

initiating, in real time, a process responsive to an event, the process including the steps of

publishing to a central repository one or more messages prompted by that event containing information from that event, respective information from the plurality of events being aggregated in the central repository where the aggregated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise,

updating the aggregated information with information from the published messages,

creating an enriched message by enriching the messages with information from that event and/or corresponding information extracted from the central repository, and

subscribing the enriched messages.

16. (Previously Presented) A ZLE virtual hub for enriched publish and subscribe operations associated with business transactions conducted by an enterprise running as a zero latency enterprise (ZLE), the system being implemented in a ZLE framework and comprising:

- one or more applications via which the business transactions are conducted; and
- an operational data store (ODS), the ODS being operatively communicative with the one or more applications such that the applications are capable to publish messages to and subscribe to messages from the ODS, the ODS being configured

- to operate as a dynamic central repository that consolidates information from across the enterprise and supports business transactional access to real time information from any of the one or more applications,

- to know what particular information any one of the applications needs in order to accomplish its task, the particular information enriching messages to which the applications subscribe, and

- to update the consolidated information with information from messages published by the applications.

17. (Original) A system as in claim 16, wherein the ODS is configured with a cluster-aware relational database management (RDBMS) functionality that is capable of handling

- periodic queries,

- message queuing and store

- state engine operations, and

- handling transactions, including insertion, updating and deletion of transactions.

18. (Original) A system as in claim 16, wherein the consolidated information can, in real-time, be accessible and available for extraction and analysis from across the enterprise, the ODS providing for a coherent view of the consolidated information, in real time, from across the enterprise.

19. (Original) A system as in claim 16, wherein the particular information for enriching messages subscribed to by an application can be information previously published by another application.

20. (Original) A system as in claim 16, further comprising:

a ZLE enriched publish-subscribe service provider module that is associated with the ODS and to which each one of applications can register with its predefined schema, the schema identifying which of a plurality of events associated with the business transactions and which types of data changes its respective application is interested in, the schema further identifying the particular information needed by its respective application.

21. (Original) A system as in claim 16, wherein for a particular number (N) of the applications, a combined number of the published and subscribed messages can be reduced from twice that particular number (2N) to a total number of 4 or 5 messages.

22. (Previously Amended) A system as in claim 16, further comprising:

an inference-based rules engine that finds an appropriate business rule, regardless of the complexity of rules or the size of any rules set, the inference-based rules engine facilitating a rules service that integrates the rules and policies of the enterprise in the ODS; and

a process-flow engine that manages a flow of the business transactions, processes, and messages between the applications integrated via the ZLE framework.

23. (Previously Presented) A system as in claim 16, wherein updating the consolidated information comprises;

providing personalized feedback and/or customized offers in real time to a customer while a customer is still engaged.

24. (Previously Presented) A method as in claim 1, where the corresponding information is mined and/or analyzed in real time.

25. (Previously Presented) A system as in claim 16, wherein the ZLE virtual hub includes a plurality of applications for performing information mining and/or analysis.

26. (Previously Presented) A system as in claim 16, wherein the ZLE virtual hub includes a plurality of technology adapters for loosely coupling applications to and logically arranging applications around the ZLE virtual hub.

27. (Previously Presented) A system as in claim 26, wherein a plurality of applications for performing information mining and/or analysis, are loosely coupled to the ZLE virtual hub via said technology adapters.

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

None.